Combinatorics, Probability and Computing (19XX) 00, 000–000. © 19XX Cambridge University Press
DOI: 10.1017/S0000000000000000 Printed in the United Kingdom

arXiv:1402.5400v1 [math.CO] 21 Feb 2014

Simple containers for simple hypergraphs

DAVID SAXTON† and ANDREW THOMASON

IMPA, Estrada Dona Castorina 110, Rio de Janeiro, Brasil 22460-320
DPMMS, Centre for Mathematical Sciences, Wilberforce Road, Cambridge, UK
saxton@impa.br and A.G.Thomason@dpmms.cam.ac.uk

We give an easy method for constructing containers for simple hypergraphs. Some applications are given; in particular, a very transparent calculation is offered for the number of H-free hypergraphs, where H is some fixed uniform hypergraph.

1. Introduction

The notion of a collection of containers for a hypergraph was introduced by the authors in [14]. A collection of containers for a hypergraph G is a collection $\mathcal{C}$ of subsets of $V(G)$ such that every independent set I is a subset of some member $C \in \mathcal{C}$. (A subset of $V(G)$ is independent if it contains no edge.)

The notion was developed further in [15] and several applications given; related methods and results were proved by Balogh, Morris and Samotij [1]. These results have since been applied by other authors.

Our purpose here is to revisit the method of [14], and to combine it with a twist that makes it much more widely applicable. It is true that the method of [15] is not too complicated, and the consequences are often best possible, but it is subtle. The method of [14], on the other hand, is not optimal; nevertheless it is very simple, and it is particularly transparent. It is sufficient, for example, for counting the number of H-free hypergraphs (see Corollary 3.2), and hence it offers a very elementary and straightforward proof of this result.

The method of [14] applies to *simple* or *linear* hypergraphs, that is, hypergraphs in which no two edges share more than one vertex. The container theorem there was as follows. We use the term *r-graph* to mean an r-uniform hypergraph, where $r \geq 2$ always.

† Supported in part by the Engineering and Physical Sciences Research Council, and CNPq bolsa PDJ.

Proposition 1.1 ([14]). *Let G be a d-regular simple r-graph. If d is large, there is a collection of sets $\mathcal{C}$ of subsets of $V(G)$ satisfying*

- *if $I \subset V(G)$ is independent, there is some $C \in \mathcal{C}$ with $I \subset C$,*
- *$|C| \leq (1 - 1/4r^2)|G|$ for every $C \in \mathcal{C}$,*
- *$|\mathcal{C}| \leq 2^{\alpha|G|}$ where $\alpha = (1/d)^{1/(2r-1)}$.*

This proposition is not quite as stated in [14], but it is pretty much explicit in the proof of Theorem 1.1 that follows Theorem 3.1.

Two drawbacks limit the applicability of Proposition 1.1. The first is that many popular container applications require containers with $e(G[C])$ small, rather than $|C|$ small. The second is that it applies only to regular r-graphs. Container results are more useful when they can be applied iteratively. That is, given an independent set I in a hypergraph G, we can apply the proposition once to obtain a container C for I, but we would then like to apply the proposition again, this time to the hypergraph $G[C]$ (of which I is still an independent set), thus finding a smaller container $C' \subset C$. If possible we would then repeat this procedure until very small containers are obtained. The snag with this procedure as it stands is that it is unlikely that $G[C]$ is regular even if G is, and so iteration is not possible. Of course, Proposition 1.1 still applies to graphs that are "somewhat" regular (indeed, this follows directly from Theorem 3.1 of [14]), but not in a particularly strong way.

Both these drawbacks can be overcome by adapting the proof of Theorem 3.1 in [14] to use the notion of *degree measure*, which we discuss in §2. This yields a version of Proposition 1.1 in which C is bounded in degree measure, namely Theorem 2.2. Iterated applications of this theorem give the following result, much stronger and more useful than Proposition 1.1, and the main theorem of the present paper.

Theorem 1.2. *Let G be a simple r-graph of average degree d. Let $0 < \delta < 1$. If d is large enough, then there is a collection of sets $\mathcal{C}$ of subsets of $V(G)$ satisfying*

- *if $I \subset V(G)$ is independent, there is some $C \in \mathcal{C}$ with $I \subset C$,*
- *$e(G[C]) < \delta e(G)$ for every $C \in \mathcal{C}$,*
- *$|\mathcal{C}| \leq 2^{\beta|G|}$ where $\beta = (1/d)^{1/(2r-1)}$.*

Observe that Theorem 1.2 differs from Proposition 1.1 only in that the condition of d-regularity is replaced by that of average degree d, and the conclusion giving a bound on $|C|$ is replaced by a bound on $e(G[C])$. This bound on $e(G[C])$ implies, for regular G, that $|C| \leq (1 - 1/r + \delta/r)n$ (see the discussion in §2), which is essentially best possible, but is nonetheless a weaker condition than the bound on the number of edges. Thus Theorem 1.2 is a generalization of Proposition 1.1.

2. Degree measure

The notion of degree measure was introduced in [15]. In the definition below, $d(v)$ is the degree of the vertex v.

Definition 2.1. *Let G be an r-graph of order n and average degree d. Let $S \subset V(G)$. The* degree measure $\mu(S)$ *of S is defined by*

$$\mu(S) \,=\, \frac{1}{nd} \sum_{v \in S} d(v)\,.$$

We note some immediate properties of degree measure. First, for any $S \subset V(G)$,

$$e(G[S]) \,\leq\, \frac{1}{r} \sum_{v \in S} d(v) \,=\, \frac{d|G|}{r} \mu(S) \,=\, \mu(S) e(G)\,, \tag{1}$$

so $\mu(S)$ small implies $e(G[S])$ small.

Moreover, sets of large measure must contain many edges. Indeed, writing $\overline{S}$ for $V(G) - S$ and $e(\overline{S}, S)$ for the number of edges meeting both $\overline{S}$ and S, we have

$$(r-1)nd\mu(\overline{S}) = (r-1) \sum_{v \notin S} d(v) \,\geq\, (r-1) e(\overline{S}, S) \,\geq\, \left\{ \sum_{v \in S} d(v) - re(G[S]) \right\},$$

that is, $(r-1)\mu(\overline{S}) \geq \mu(S) - re(G[S])/nd$. Since $\mu(\overline{S}) = 1 - \mu(S)$ this means

$$e(G[S]) \,\geq\, (\mu(S) - 1 + \frac{1}{r})\, nd\,. \tag{2}$$

In particular, the measure of an independent set cannot exceed $1 - 1/r$. Furthermore, if G is regular, then degree and uniform measures coincide; in this case, the inequality $e(G[S]) \leq \delta e(G)$ together with (2) implies $|S| \leq (1 - 1/r + \delta/r)n$, as mentioned in the introduction.

We can now state the theorem which is at the heart of the present paper. This theorem is already sufficiently powerful for obtaining non-trivial results, such as in list colouring.

Theorem 2.2. *Let G be a simple r-graph of average degree d. If d is large, there is a collection of sets $\mathcal{C}$ of subsets of $V(G)$ satisfying*

- *if $I \subset V(G)$ is independent, there is some $C \in \mathcal{C}$ with $I \subset C$,*
- *$\mu(C) \leq 1 - 1/4r^2$ for every $C \in \mathcal{C}$,*
- *$|\mathcal{C}| \leq 2^{\alpha |G|}$ where $\alpha = (1/d)^{1/(2r-1)}$.*

The proof of Theorem 2.2 follows quite closely the proof of Theorem 3.1 in [14], but modifications are needed to accommodate the presence of both uniform and degree measures. However, nothing stronger than Markov's inequality is needed.

The spirit of the proof is readily explained. We need to identify a set of vertices that are not in I; then C will be the remaining vertices. We shall show that there are three small subsets R, S and T of $V = V(G)$, such that R, S and T determine such a set $V \setminus C$ disjoint from I. This means that the number of different container sets C that are so specified is at most the number of triples of small subsets (R, S, T); this number is not large and this is where the bound on $|\mathcal{C}|$ comes from.

How, then, can we specify R, S and T in such a way as to enable us to identify a set $V \setminus C$ of vertices not in I? There are no edges with all r vertices inside I, but there are

many edges altogether. So there must be a number j, $0 \le j < r$, such that there are significantly fewer edges with $j+1$ vertices in I than there are edges with j vertices in I. We might then expect to find a substantial set $D \subset V \setminus I$ of vertices each lying in many of the latter kind of edges. So we pick small subsets $R \subset I$ and $S \subset V \setminus I$ at random, and look at the set

$$\Gamma_j(R,S) = \{v \in V : \text{ there is an edge } \{v\} \cup f \cup g \text{ with } f \in R^{(j)} \text{ and } g \in S^{(r-j-1)}\},$$

where $R^{(j)} = \{Y \subset R : |Y| = j\}$, etc. Notice that $\Gamma_j(R,S)$ is determined by R and S. If we write $T = \Gamma_j(R,S) \cap I$ then clearly $C = (V \setminus \Gamma_j(R,S)) \cup T$ is a container for I that is specified by (R,S,T). Now R and S are small by definition, and we expect T also to be small, because there are few edges with $j+1$ vertices in I. On the other hand, vertices of D have a good chance of lying inside $\Gamma_j(R,S)$, so we expect $\Gamma_j(R,S)$ to contain much of D and so have substantial measure, meaning that $\mu(C)$ is bounded away from one. This is the heart of the proof.

Proof of Theorem 2.2. Let $V = V(G)$ be the vertex set of G of size $n = |V|$ and $E = E(G)$ the edge set. For sets $R, S \subset V$ and $0 \le j \le r-1$, let $\Gamma_j(R,S)$ be as defined above. Given subsets $R, S, T \subset V$, let

$$C_j(R,S,T) = \begin{cases} V \setminus (\Gamma_j(R,S) \setminus T) & \text{if } \mu(\Gamma_j(R,S) \setminus T) \ge 1/4r^2 \\ \emptyset & \text{otherwise.} \end{cases}$$

Note that $\mu(C_j(R,S,T)) \le 1 - 1/4r^2$ by definition. We will show that for every independent set I, there are small subsets $R, S, T \subset V$ such that $I \subset C_j(R,S,T)$. Specifically, let

$$u = \frac{1}{\sqrt{3r}} \left(\frac{6r}{d}\right)^{1/2(r-1)} \quad \text{and} \quad q = 15ru\,.$$

Note that q is small if d is large (depending on r). We now define the collection $\mathcal{C}$ by $\mathcal{C} = \{C_j(R,S,T) : 0 \le j \le r-1,\ |R|, |S|, |T| \le qn\}$. Then

$$|\mathcal{C}| \ \le\ r(qn)^3 \binom{n}{qn}^3 \ \le\ r(qn)^3 \left(\frac{ne}{qn}\right)^{3qn} \ \le\ 2^{\alpha n}$$

for d sufficiently large, where $\alpha = (1/d)^{1/(2r-1)}$. This collection $\mathcal{C}$ will satisfy the conditions of the lemma.

Fix an independent set I. For a subset $A \subset V$ with $I \subset A$, and for $0 \le j \le r$, we define the set of edges

$$E_j(A) = \{e \in E : e \subset A,\ |e \cap I| \ge j\}\,.$$

Let $P(j)$ be the statement

$$\text{for all } A \subset V \text{ with } I \subset A \text{ and } \mu(A) \ge 1 - 1/2r + j/2r^2,\ |E_j(A)| \ge ndu^j/2r \text{ holds}\,.$$

Statement $P(0)$ is true by (2), since $|E_0(A)| = e(G[A])$. Statement $P(r)$ is false, because I is independent and so $E_r(A) = \emptyset$. There must therefore exist $j \in \{0, 1, \ldots, r-1\}$ such that $P(j)$ is true and $P(j+1)$ is false. Fix a set A witnessing the falsity of $P(j+1)$; thus

$I \subset A$, $\mu(A) \geq 1 - 1/2r + (j+1)/2r^2$ and $|E_{j+1}(A)| < ndu^{j+1}/2r$. For $v \in A$, let

$$F_j(v) = \{e \in E : v \in e,\, e \in E_j(A),\, e \notin E_{j+1}(A)\} = \{e : v \in e \subset A,\, |e \cap I| = j\}\,.$$

Let $D = \{v \in A \setminus I : |F_j(v)| \geq du^j(1-u)/2r\}$. Note that $I \subset A \setminus D$.

Consider an edge $e \in E_j(A \setminus D)$ with $e \notin E_{j+1}(A)$. Then $e \subset A \setminus D$ and $|e \cap I| = j$. Since $j < r$ we can pick $v \in e$ with $v \notin I$. Now $E_j(A \setminus D) \subset E_j(A)$ so, by definition of $F_j(v)$, we have $e \in F_j(v)$. Moreover, since $v \notin I$, the definition of D and the fact that $v \notin D$ imply $|F_j(v)| < du^j(1-u)/2r$. Therefore, the total number of edges in $E_j(A \setminus D)$ but not in $E_{j+1}(A)$ is less than $|A \setminus D|du^j(1-u)/2r \leq ndu^j(1-u)/2r$. By the choice of A as witness set, we know that $|E_{j+1}(A)| < ndu^{j+1}/2r$ and so $E_j(A \setminus D) < ndu^j(1-u)/2r + ndu^{j+1}/2r = ndu^j/2r$. Since $P(j)$ is true, this means $\mu(A \setminus D) < 1 - 1/2r + j/2r^2$. But $\mu(A) \geq 1 - 1/2r + (j+1)/2r^2$ and therefore $\mu(D) > 1/2r^2$.

Let $p = (6r/du^j)^{1/(r-1)}$, so $p^{r-1}du^j = 6r$. Since $j \leq r-1$, we observe that

$$p \;\leq\; \left(\frac{6r}{d}\right)^{1/(r-1)} \frac{1}{u} \;=\; \sqrt{3r}\left(\frac{6r}{d}\right)^{1/2(r-1)} \;=\; 3ru \;=\; \frac{q}{5}\,.$$

Let $R \subset I$ and $S \subset A \setminus I$ be random sets where each vertex (of I and $A \setminus I$ respectively) is included independently with probability p. By Markov's inequality, the inequalities $|R| \leq 5pn \leq qn$ and $|S| \leq 5pn \leq qn$ each hold with probability at least $4/5$. Let $T = \Gamma_j(R,S) \cap I$. Then clearly, $I \subset C_j(R,S,T)$ provided $\mu(\Gamma_j(R,S) \setminus T) \geq 1/4r^2$. So to complete the proof, it is enough to show that the inequalities $|T| \leq qn$ and $\mu(\Gamma_j(R,S) \setminus T) \geq 1/4r^2$ each hold with probability at least $4/5$, because then, with positive probability, all four inequalities $|R|, |S|, |T| \leq qn$ and $\mu(\Gamma_j(R,S) \setminus T) \geq 1/4r^2$ will hold.

A vertex $v \in I$ will be included in $\Gamma_j(R,S)$ (i.e., in T) if it lies in an edge e with $e = \{v\} \cup f \cup g$, $f \in R^{(j)}$, $g \in S^{(r-j-1)}$. Therefore $e \subset A$ and $|e \cap I| = j+1$, which means $e \in E_{j+1}(A)$. We know $|E_{j+1}(A)| < ndu^{j+1}/2r$. For an edge $e \in E_{j+1}(A)$ with $|e \cap I| = j+1$, there are $j+1$ partitions of e of the form $e = \{v\} \cup f \cup g$ with $v \in I$, $f \in I^{(j)}$ and $g \in (A \setminus I)^{(r-j-1)}$. For each such partition, the probability that both $f \in R^{(j)}$ and $g \in S^{(r-j-1)}$ is p^{r-1}. So the expected size of T is at most

$$rp^{r-1}ndu^{j+1}/2r \;=\; 3run \;=\; qn/5.$$

Applying Markov's inequality again implies that $|T| \leq qn$ with probability at least $4/5$.

Recall that $D \cap I = \emptyset$ by definition of D, and so in particular $D \cap T = \emptyset$. Let $D^* = D \setminus \Gamma_j(R,S)$. Then $D \setminus D^* \subset \Gamma_j(R,S) \setminus T$, and so $\mu(\Gamma_j(R,S) \setminus T) \geq \mu(D \setminus D^*) = \mu(D) - \mu(D^*)$. Let $v \in D$. Then $|F_j(v)| \geq du^j(1-u)/2r > 2du^j/5r$ (since u is small). Each $e \in F_j(v)$ has a partition $e = f \cup g$ with $f \in I^{(j)}$ and $g \in (A \setminus I)^{(r-j)}$, where $v \in g$ because $v \notin I$. The probability that $f \subset R$ and $g - \{v\} \subset S$ is p^{r-1} and, because G is simple, these events over all $e \in F_j(v)$ are independent. Hence the probability that $v \in D^*$, that is, $v \notin \Gamma_j(R,S)$, is at most

$$(1-p^{r-1})^{|F_j(v)|} \leq \exp\{-2p^{r-1}du^j/5r\} = \exp\{-12/5\} < 1/10.$$

Now $\mu(D^*) = (1/nd)\sum_{v \in D^*} d(v) = (1/nd)\sum_{v \in D} d(v) I_v$ where I_v is the indicator of the event $v \in D^*$. Taking expectations, $\mathbb{E}\mu(D^*) = (1/nd)\sum_{v \in D} d(v)\mathbb{E}(I_v) < \mu(D)/10$, since $\mathbb{E}(I_v) = \Pr(v \in D^*) < 1/10$. Markov's inequality implies that, with probability at

least $4/5$, $\mu(D^*) \le \mu(D)/2$ holds, and hence $\mu(\Gamma_j(R,S)\setminus T) \ge \mu(D) - \mu(D^*) \ge \mu(D)/2 > 1/4r^2$. This completes the proof. □

We remark that the proof shows the theorem to be true for a smaller value of α, namely $c(r)(\log d)/d^{1/2(r-1)}$ for some function $c(r)$ of r, but the results of [15] are better still, with $\alpha = c(r)(\log d)/d^{1/(r-1)}$, so we keep the present value for simplicity. One might wonder why the bound here on $|\mathcal{C}|$ is worse than the bound in [15]. It is not because of the random choice of R and S, because in the context of the present algorithm random choice is quite efficient, and a deterministic choice is unlikely to be much better. The reason that the present method is relatively inefficient is that it uses edges with exactly j vertices in I for one value of j only, and ignores all other edges. The methods of [15] and [1], which are unrelated to the present method, are not lengthy to describe but are nonetheless crafted carefully to use all edges and to be as efficient as possible.

Applying Theorem 2.2 repeatedly, as described earlier, we obtain the main theorem.

Proof of Theorem 1.2. As we remarked earlier, let us apply Theorem 2.2 to G itself, and then again to each container so obtained, then to each of the new containers, and so on for each container with at least $\delta e(G)$ edges, until we obtain a collection $\mathcal{C}$ of containers C with $e(G[C]) < \delta e(G)$. Since, by (1), each application of Theorem 2.2 decreases the fraction of edges by $1-1/4r^2$, $\mathcal{C}$ is obtained after at most $k = \lceil (\log \delta)/\log(1-1/4r^2)\rceil + 1$ levels of iteration, and so $|\mathcal{C}| \le 2^{k\alpha|G|}$, where α is the maximum over all applications of Theorem 2.2. If $e(G[C]) \ge \delta e(G)$ then the average degree of $G[C]$ is at least δd, and the result follows, provided, for the sake of a clean statement, the reader will indulge us by taking the value $\alpha(d) = c(r)(\log d)/d^{1/2(r-1)}$ rather than the weaker bound explicit in Theorem 2.2. □

3. Applications

As remarked earlier, for regular hypergraphs, the condition $e(G[C]) < \delta e(G)$ implies $|C| < (1 - 1/r + \delta/r)n$. Plugging this value into Theorem 2.1 of [14] immediately improves by a factor of two the bound on the list colouring number in Theorem 1.1 of [14]; nevertheless the bound obtained remains a factor of two worse than the bound in Theorem 1.3 of [15], which is probably best possible.

We now give an application of Theorem 1.2 in a situation where the hypergraph of interest is not simple. In what follows, H is a fixed ℓ-graph. We call another ℓ-graph *H-free* if it has no subgraph isomorphic to H. The maximum size of an H-free ℓ-graph on N vertices is denoted by $\mathrm{ex}(N,H)$, and $\pi(H) = \lim_{N\to\infty} \mathrm{ex}(N,H)\binom{N}{\ell}^{-1}$ is the limiting maximum density of H-free ℓ-graphs.

Theorem 3.1. *Let H be an ℓ-graph and let $\epsilon > 0$. Then, if N is large enough, there exists a collection $\mathcal{C}$ of ℓ-graphs on vertex set $[N]$ such that*

- *every H-free ℓ-graph on vertex set $[N]$ is a subgraph of some $C \in \mathcal{C}$,*
- *every $C \in \mathcal{C}$ has at most $\epsilon N^{v(H)}$ copies of H, and $e(C) \le (\pi(H)+\epsilon)\binom{N}{\ell}$,*
- *$\log|\mathcal{C}| \le N^{\ell-\sigma}$ where $\sigma = 1/2e(H)$.*

The meaning of this theorem is that every H-free ℓ-graph is a subgraph of one of just a few ℓ-graphs that are nearly H-free. The strength of the result can be measured by the bound on $\log|\mathcal{C}|$. The bound N^ℓ is, of course, trivial, but any bound where σ is some positive constant is worthwhile. It is not possible for σ to exceed the value $m(H) = \max_{H'\subset H,\, e(H')>1} (e(H')-1)/(v(H')-\ell)$, and in fact a best possible bound was obtained in [15, Theorem 1.3], but the method here is simpler. Any value of $\sigma > 0$, such as that given by Theorem 3.1, immediately gives the following corollary, because each graph C in the statement of the theorem has at most $2^{(\pi(H)+\epsilon)\binom{N}{\ell}}$ subgraphs.

Corollary 3.2. *Let H be an ℓ-graph. The number of H-free ℓ-graphs on vertex set $[N]$ is $2^{(\pi(H)+o(1))\binom{N}{\ell}}$.*

This corollary is the same as [15, Corollary 1.4]. In the case $\ell = 2$, this corollary was proved for complete H by Erdős, Kleitman and Rothschild [6] and for general H by Erdős, Frankl and Rödl [5]. Nagle, Rödl and Schacht [12] proved it for general ℓ using hypergraph regularity methods. The present paper offers the simplest known proof.

Theorem 1.2 can, in a similar way, be used to give a simple way to count the number of ℓ-graphs which have no *induced* copy of H, and more generally to evaluate the probability that a random ℓ-uniform hypergraph $G^{(\ell)}(n,p)$ contains no induced copy of H. For $\ell = 2$, the value when $p = 1/2$ was determined by Prömel and Steger [13] and for general p by Bollobás and Thomason [2] (see also Marchant and Thomason [10]). For general ℓ the value for $p = 1/2$ was given by Dotson and Nagle [4], again using hypergraph regularity techniques. We don't give details of the result, which is identical to [15, Theorem 2.5]. We merely point out that it can be derived from a container theorem, as demonstrated in [15], and that the container theorem presented here can be used instead, via an argument very similar to the one used to prove Theorem 3.1.

Another application of Theorem 3.1 is the following "sparse Turán theorem". Here the value of σ does affect the strength of the application.

Corollary 3.3. *Let H be an ℓ-graph and let $0 < \gamma < 1$. For some $c > 0$, for N sufficiently large and for $p \ge cN^{-\sigma}$, where $\sigma = 1/2e(H)$, the following event holds with probability greater than $1 - \exp\{-\gamma^3 p\binom{N}{\ell}/512\}$:*

every H-free subgraph of $G^{(\ell)}(N,p)$ has at most $(\pi(H)+\gamma)p\binom{N}{\ell}$ edges.

A stronger version of this corollary, with $\sigma = 1/m(H)$, was conjectured by Kohayakawa, Łuczak and Rödl [8]; it was proved in the case of strictly balanced H by Conlon and Gowers [3] and in full generality by Schacht [16]. The strong version follows easily from [15, Theorem 1.3], as shown in [15], and the same argument gives Corollary 3.3 from Theorem 3.1, so we do not give details here. We remark that the point of the corollary is how small the value of p can be made: Szemerédi's regularity lemma allows $p = o(1)$. We note that Kohayakawa, Rödl and Schacht [9] and Szabó and Vu [17] both proved the corollary for complete 2-graphs with $\sigma = 1/(v(H)-1)$ (slightly better in the case of [17]), but

again we believe the present proof is the shortest for some $\sigma > 0$. It yields in a similar fashion weak versions of the other so-called KŁR conjectures.

The proof of Theorem 3.1 consists of finding a set of containers for the independent sets in the hypergraph $G = G(N, H)$, which is defined as follows. The $n = \binom{N}{\ell}$ vertices of G are the ℓ-sets in $[N]$, that is, $V(G) = [N]^{(\ell)}$. The edges of G are the subsets of size $e(H)$ of $V(G)$ that form an ℓ-graph isomorphic to H.

Given a subset $S \subset V(G)$, we can regard S as the edges of an ℓ-graph with vertex set $[N]$. The subset S is independent in G if and only if S, regarded as an ℓ-graph, is H-free. A set of containers $\mathcal{C}$ for the independent sets in G is thus a set of ℓ-graphs on vertex set N such that every H-free graph is a subset of one of these container graphs. Thus Theorem 3.1 is a statement about the existence of a collection $\mathcal{C}$ of containers for $G(N, H)$ having certain properties.

The stronger [15, Theorem 1.3] was obtained by applying a container result directly to $G(N, H)$. Here, we use the simpler Theorem 1.2 to give a set of containers with slightly weaker properties. We cannot apply Theorem 1.2 directly to $G(N, H)$ because this hypergraph is not simple. We therefore apply it instead to a subgraph $G_{\text{simple}}(N, H)$ of $G(N, H)$. Each independent set of $G(N, H)$ is independent in $G_{\text{simple}}(N, H)$, so containers for $G_{\text{simple}}(N, H)$ will also be containers for $G(N, H)$. To show that these containers have the properties claimed in Theorem 3.1, we need the following lemma.

Lemma 3.4. *Let $\eta > 0$ and $0 < \rho < 1$. Then, if N is large enough, there exists a simple sub-hypergraph $G_{\text{simple}} = G_{\text{simple}}(N, H)$ of $G = G(N, H)$, such that $V(G_{\text{simple}}) = V(G)$ and G_{simple} has average degree at least N^ρ. Moreover, for all $S \subset V(G)$, if $e(G[S]) \geq \eta e(G)$, then $e(G_{\text{simple}}[S]) \geq \eta e(G_{\text{simple}})/2$.*

Given this lemma, the proof of Theorem 3.1 follows at once, using the supersaturation theorem of Erdős and Simonovits [7], which itself has a very straightforward proof.

Proposition 3.5 (Erdős and Simonovits [7]). *Let H be an ℓ-graph and let $\epsilon > 0$. There exists N_0 and $\eta > 0$ such that if C is an ℓ-graph on $N \geq N_0$ vertices containing at most $\eta N^{v(H)}$ copies of H then $e(C) \leq (\pi(H) + \epsilon)\binom{N}{\ell}$.*

Proof of Theorem 3.1. Let $\epsilon > 0$ be as given in the conditions of the theorem. Then let $\eta > 0$ be given by Proposition 3.5. We may of course assume that $\eta \leq \epsilon$. Choose $\rho < 1$ so that $\rho/(2e(H)-1) > 1/2e(H)$. Then apply Lemma 3.4 to obtain G_{simple}. Apply Theorem 1.2 to G_{simple} with $\delta = \eta/2$, with $n = \binom{N}{\ell}$ and $d \geq N^\rho$, so d is large if N is large, to obtain a collection $\mathcal{C}$ for the independent sets in G_{simple}. As remarked before, every H-free ℓ-graph I on vertex set $[N]$ is an independent set in G_{simple} and is therefore contained in some subset $C \in \mathcal{C}$, which itself can be regarded as an ℓ-graph on vertex set $[N]$. We have $|\mathcal{C}| \leq 2^{\beta n}$ where $\beta = (1/d)^{1/(2e(H)-1)}$. Since $\rho/(2e(H) - 1) > 1/2e(H)$ we have $\log|\mathcal{C}| \leq N^{\ell - 1/2e(H)}$, as claimed.

All that remains, then, is to verify the second assertion of the theorem. In the assertion, the number of copies of H in C is the same as $e(G[C])$. By Theorem 1.2, $e(G_{\text{simple}}[C]) < \delta e(G_{\text{simple}})$. Since $\delta = \eta/2$, Lemma 3.4 shows $e(G[C]) < \eta e(G) \leq \epsilon e(G)$. Now $e(G)$ is

the number of copies of H with vertices in $[N]$ and so $e(G) < N^{v(H)}$. So Proposition 3.5 implies $e(C) \leq (\pi(H)+\epsilon)\binom{N}{\ell}$, completing the proof. □

Proof of Lemma 3.4. We form G_{simple} by randomly choosing edges of G and then deleting a few so that the result is a simple hypergraph. Observe that $\binom{N}{v(H)} \leq e(G) \leq N^{v(H)}$, so that $e(G) = \Theta(N^h)$ where $h = v(H)$. We may assume that H has more than one edge, and so $h \geq \ell + 1 \geq 3$. Call a pair e, e' of edges of G with $|e \cap e'| \geq 2$, an *overlapping pair*. Notice that the number of overlapping pairs is the number of copies H, H' of H with vertices in $[N]$ that have at least two ℓ-edges in common: H and H' must share at least $\ell+1$ vertices and so the number of overlapping pairs is $O(N^{2h-\ell-1})$.

Pick a number ρ' with $\rho < \rho' < 1$. Let G' be a subgraph of G formed by picking edges independently and at random with probability $p = N^{-h+\ell+\rho'}$. Let E be the number of edges of G'. We make use of standard bounds on the tail of the binomial distribution, to wit, if $X \sim \text{Bi}(m,p)$ then $\Pr\{X \leq (3/4)\mathbb{E}X\} \leq e^{-\mathbb{E}X/40}$, and the same bound holds for $\Pr\{X \geq (5/4)\mathbb{E}X\}$ (see for example [11, Corollary 2.3]). So if A is the event $\{3\mathbb{E}E/4 \leq E \leq 5\mathbb{E}E/4\}$ then A holds with high probability, certainly more than $2/3$. Observe that if A holds then G' has $\Theta(N^{\ell+\rho'})$ edges.

Let F be the number of overlapping pairs in G'. Then $\mathbb{E}F = O(p^2 N^{2h-\ell-1}) = O(N^{\ell+2\rho'-1}) = o(N^{\ell+\rho'})$. Let B be the event $\{F \leq 3\mathbb{E}F\}$. By Markov's inequality, B holds with probability at least $2/3$.

For each $S \subset V(G)$, let C_S be the event that both $e(G[S]) \geq \eta e(G)$ and $e(G'[S]) \leq 3pe(G[S])/4$ hold. Then the probability that C_S occurs is at most $\exp(-pe(G[S])/40) = \exp(-p\Theta(N^h)) = \exp(-\Theta(N^{\ell+\rho'}))$. Let C be the event that C_S does not hold for any $S \subset V(G)$. There are $2^{|V(G)|}$ subsets S, so the probability that C fails to hold is at most $\exp(N^\ell - \Theta(N^{\ell+\rho'})) = o(1)$.

There is therefore a positive probability that A, B and C all hold. Let G' be a graph for which they all do hold, remove an edge from each overlapping pair, and call the result G_{simple}. This graph has no overlapping pairs and so is simple. The number of edges is $E - F$. Since A and B hold, we have $E = \Theta(N^{\ell+\rho'})$ and $F = o(N^{\ell+\rho'})$, so $F = o(E)$ and $E - F = \Theta(N^{\ell+\rho'}) > N^{\ell+\rho}$. The graph has fewer than N^ℓ vertices and so its average degree exceeds N^ρ. Finally, let $S \subset V(G)$ be such that $e(G[S]) \geq \eta e(G)$. The event C holds, and so C_S does not; thus $e(G'[S]) \geq 3pe(G[S])/4 \geq 3p\eta e(G)/4 \geq (3/4)\eta(4/5)E$, the last inequality holding because A holds. Therefore $e(G_{\text{simple}}[S]) \geq 3\eta E/5 - F$. But $F = o(E)$ so $e(G_{\text{simple}}[S]) \geq 3\eta E/5 - F \geq \eta(E-F)/2 = \eta e(G_{\text{simple}})/2$, which completes the proof. □

References

[1] J. Balogh, R. Morris and W. Samotij, Independent sets in hypergraphs, (to appear in *J. Amer. Math. Soc.*).

[2] B. Bollobás and A. Thomason, The structure of hereditary properties and colourings of random graphs, *Combinatorica* **20** (2000) 173–202.

[3] D. Conlon and W.T. Gowers, Combinatorial theorems in sparse random sets, submitted.

[4] R. Dotson and B. Nagle, Hereditary properties of hypergraphs, *J. Combinatorial Theory (Ser. B)* **99** (2009), 460–473.

[5] P. Erdős, P. Frankl and V. Rödl, The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent, *Graphs Combin.* **2** (1986), 113–121.
[6] P. Erdős, D.J. Kleitman and B.L. Rothschild, Asymptotic enumeration of K_n-free graphs, *Colloquio Internazionale sulle Teorie Combinatorie (Rome, 1973)* **2** (1976), 19–27.
[7] P. Erdős and M. Simonovits, Supersaturated graphs and hypergraphs, *Combinatorica* **3** (1983), 181–192.
[8] Y. Kohayakawa, T. Łuczak and V. Rödl, On K_4-free subgraphs of random graphs, *Combinatorica* **17** (1997), 173–213.
[9] Y. Kohayakawa, V. Rödl and M. Schacht, The Turán theorem for random graphs, *Combinatorics, Probability and Computing* **13** (2004), 61–91.
[10] E. Marchant and A. Thomason, The structure of hereditary properties and 2-coloured multigraphs, *Combinatorica* **31** (2011), 85–93.
[11] S. Janson, T. Łuczak and A. Ruciński, *Random graphs* (2000), Wiley.
[12] B. Nagle, V. Rödl and M. Schacht, Extremal hypergraph problems and the regularity method, in "Topics in discrete mathematics", *Algorithms Combin.* **26** (2006), 247–278.
[13] H.-J. Prömel and A. Steger, Excluding induced subgraphs III: a general asymptotic, *Rand. Struct. Alg.* **3** (1992), 19–31.
[14] D. Saxton and A. Thomason, List colourings of regular hypergraphs, *Combinatorics, Probability and Computing* (2011).
[15] D. Saxton and A. Thomason, Hypergraph containers, arXiv:1204.6595 [math.CO].
[16] M. Schacht, Extremal results for random discrete structures, submitted.
[17] T. Szabó and V.H. Vu, Turán's theorem in sparse random graphs, *Random Structures and Algorithms* **23** (2003), 225–234.